

534-63
12P.

185116

Telerobot Task Planning and Reasoning: Introduction to JPL Artificial Intelligence Research

D.J. Atkinson
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109

JJ574450

ALL right

A. Introduction

The goal of building robots with increasing amounts of autonomy from direct human control or supervision is a long term goal of the Telerobot project. Among the primary missions for such robots are tasks in the areas of on-orbit satellite servicing, inspection, and various assembly operations associated with the space station. A research and development program with a series of demonstrations of increasing robot autonomy has been planned. This paper discusses some of the technical, methodological, and logistical issues of producing the artificial intelligence capabilities required by the long-term (1993 through 2000-era) demonstrations of these robot systems. A substantial research effort must begin now if the demonstration objectives are to be met on schedule. It is no coincidence that many of the same objectives for Telerobot artificial intelligence research are common to both the Telerobot and System Autonomy programs. Many of the capabilities and associated research required for autonomy transcend the particular application.

The planned series of Telerobot demonstrations provides a relatively clear progression of increasing sophistication in the artificial intelligence capabilities required. As a directed research and development program, these demonstrations provide considerable *applications pull* in the kind of research issues which may be addressed. However, applications pull alone can easily become overly restrictive to new research ideas and also unrealistic about existing or projected technological capability. Basic research, on the other hand, provides the *technology push* towards new ideas. Ultimately, this push results in demonstrations and application of more powerful capabilities. Technology push by itself runs the risk of being distractable and undirected towards economically realistic and desirable application goals. Clearly, a mixture of both applications pull and technology push is necessary. Achieving a sensible balance between these two opposing forces will be major criteria of on-going and overall success of the Telerobot project.

Due to the scope of the artificial intelligence technology required by Telerobot, one of the conclusions to draw from this document is that our goals for Telerobot artificial intelligence will be difficult to achieve unless we make maximal use of all available research resources. NASA probably does not have the desire or capability to fund all of the artificial intelligence research which will be required by

Telerobot, nor should NASA duplicate research conducted under other auspices. We will need to draw extensively upon research results generated by the Systems Autonomy Program at the NASA Ames Research Center laboratory, and externally funded laboratories at major universities and within the AI industry.

B. THE LONG-RANGE GOAL: CAPABILITIES

This section details the specific capabilities which will be required in the area of Artificial Intelligence for the long-term Telerobot demonstrations leading to operational 2000-era autonomous robots. These capabilities should be regarded as *necessary* for the specific objectives of the demonstration program. However, they should not be construed to be a *sufficient* set of capabilities. This can only be established by continuing research, experimentation, and the demonstrations themselves. Subsequent sections will provide detail on some of the technical issues which must be addressed.

B.1 TASK PLANNING AND SCHEDULING

A number of processes must operate at the **Task Level** to robustly control the achievement of goals. One of these is task planning and scheduling. Its function is to select and schedule activities whose collective effects logically achieve the desired goals. In the Telerobot, some of the planning and scheduling capabilities required are:

- **Use of highly structured procedures.**
Many of the tasks which the robot will carry out have been engineered beforehand, e.g., there are satellite servicing procedures which must be followed in a definite, prescribed sequence. The robot should have the ability to select the appropriate procedure and combine it with other elements of its activity plan. It will be a rare case when predefined procedures can be simply "invoked" without regard to plan interactions or the potential for error due to real-world uncertainty.
- **Creation of ad hoc routines.**
The robot must be able to reason about the task requirements of novel goals. All situations in which the robot is expected to behave cannot be conceived or engineered beforehand. There will undoubtedly be circumstances when a structured procedure has not been developed. This will certainly be the case when a structured procedure fails and error recovery processes must create appropriate plans to handle the novel circumstances. The robot must be able to freely mix planning using structured procedures and ad hoc routines.
- **Management of task resources.**
In the accomplishment of any goal, the robot can draw only upon limited resources. There are a finite number of effectors available, a finite number of sensors to employ, a finite number of interchangeable or unique tools or replacement parts, etc. Management of resource constraints is a fundamental component of creating viable plans.

- **Planning with uncertainty.**

The ability to formulate robust plans when there is considerable uncertainty about the run-time environment, the robot's ability to affect it, or the knowledge used in planning (see Uncertainty management below).

- **Multi-agent cooperation.**

The ability to formulate plans which provide for specific tasks to be accomplished in a coordinated fashion by other intelligent agents, robot or human. Also, the ability to formulate plans which are robust in circumstances where the actions of other intelligent agents are not prescribed or are otherwise unpredictable.

B.2 SPATIAL PLANNING AND REASONING

The function of spatial planning and reasoning processes is to bring the logical procedures of activity planning into the real-world by making geometric, physical, or temporal constraints on activities an integral part of the task level control process. This may require a theory of manipulative processes which support planning for handling, service, repair, construction, and inspection. Some of the capabilities which are required are:

- **Reasoning about robot, workspace, and workpiece geometry.**

Geometric constraints on tasks can play a significant role in the selection and sequencing of appropriate actions by the robot. Given a device to be disassembled for repair, the selection and sequencing of manipulations is dependent on the geometric configuration of the device, e.g., you can't remove an orbital replacable unit (ORU) without removing an access panel. Contact and attachment constraints easily generate other examples where geometric reasoning is important.

- **Reasoning about physical processes.**

The ability to plan for or around the effects of physical processes is important. When a tool is released in zero-g, the robot must understand how physical processes such as inertia will affect the state of the tool, e.g., will it float away and if so, where and at what speed? Often only qualitative answers are required in the planning process (e.g., the tool will rotate, float left, etc.), but the robot must be able to utilize precise quantitative information where necessary (e.g., to direct the movement of a manipulator towards a wrench which is floating away). Another example would be reasoning about the flow of fluid, such as fuel, through a hose in a refueling operation. The planning process must take into account physical constraints on planned activities.

- **Planning for workpiece and robot positioning.**

The robot must be able to reason about how to position itself and the workpiece to enable required operations, and when this is impossible, to provide the appropriate constraints for activity planning. In zero-G, without the default

- constraints and orientation provided by Earth gravity, a mobile robot has considerable latitude in this area which should be exploited.

- **Planning for robot movement.**

The ability to plan reasonable and safe manipulator trajectories through a cluttered workspace is a critical component of task level control. Collision avoidance with fixed obstacles, stationary but movable obstacles, and moving obstacles is a significant problem.

- **Mobility.**

The ability to plan the movement of the robot and navigate in real-time through a 3-dimensional environment. The robot should have the ability to maintain its position relative to moving workspaces (i.e., station keeping) as well as plan its movements so as not to conflict with other mobile objects (e.g., other robots, astronauts on EVA, or spacecraft).

- **Temporal Reasoning.**

Reasoning about the duration and timing of activities and physical processes introduces additional constraints on the task planning process. The scheduling of activities to accommodate real-world events, including the actions of other agents, will rely on an effective temporal reasoning ability.

- **World modelling.**

The ability to maintain a coherent geometric and physical model of the world as it dynamically changes as a result of robot actions, the actions of other agents, and controlled or uncontrolled physical processes is critical. Virtually all task level control processes must have access to a (reasonably) consistent set of beliefs about the state of the world in order to make assumptions for planning, diagnosis, error recovery, and other functions. When conflicting information and ambiguity are introduced, this can become a serious problem.

B.3 EXECUTION MONITORING

In a sometimes malevolent and only partially-modelled real world, a robot executing a plan may not actually achieve the desired effects of each action while still obeying any constraints imposed by the plan. There are a host of potential problems which may create a divergence at execution time from the expected effects of actions as described in the plan. It is the role of execution monitoring processes to determine whether the plan is executing nominally, and to track the state of the world during plan execution. A number of capabilities are important:

- **Verification and sensor planning.**

The ability to determine how best to employ sensors during plan execution to verify that robot actions are achieving their intended effects, and to notice other events which may affect successful plan execution. Planning and scheduling, resource management, and sensor modelling techniques are aspects of this process.

- **Expectation generation.**

In order to determine that the intended effects of a plan are actually occurring, qualitative and quantitative predictions about what information will be obtained on sensors during plan execution must be generated.

- **Situation assessment and sensor data fusion.**

In lieu of perfect models of the sensors and the world, the expectations about information on various sensors will be approximate at best. Techniques must be developed to make accurate judgements of whether expectations are satisfied or violated from partial data in predictions and obtained from sensors. Frequently, information from multiple sensors will need to be coordinated in this process.

• **Recognition of unexpected events.**

The ability to recognize and characterize events or states of the world which are unanticipated. In the formulation of plans, the robot will not have the ability to predict all events which may affect the successful execution of the plan. In order to robustly accommodate potential conflicts at plan execution time, the robot must first be able to notice and describe unexpected situations.

B.4 DIAGNOSIS AND ERROR INTERPRETATION

When errors or unexpected events are detected during plan execution, diagnosis and other error interpretation processes must attempt to discover the source of the problem. In the absence of a precise diagnosis, error recovery or restart procedures will be coarse and may needlessly cause the duplication of work successfully completed or even the overall failure to achieve the required goals of the robot. On the other hand, a precise diagnosis may indicate a simple procedure to recover (e.g., regrip a tool if it slips during usage) or the need to plan a more substantial recovery (e.g., when the satellite's access panel is not attached as described in the CAD database). The following capabilities are required:

• **Localizing system failures.**

Identification of single and multiple point failures of the hardware of the robot or associated systems. This type of diagnosis is critical for determining the status of the robot and its ability to accomplish required tasks.

• **Knowledge-base criticism.**

Occasionally, the knowledge which the robot uses to plan actions, monitor, or otherwise solve problems will be missing or in error. The capability to spot these gaps or errors will be essential, especially when machine learning becomes an important process.

• **Debugging lower-level controllers.**

In some cases, lower-level controllers may operate incorrectly for a variety of reasons. The ability to make this decision (and possibly suggest a restart of the controller or other recovery action) is important.

• **Diagnostic test design.**

The robot should have the ability to devise and utilize diagnostic tests to discriminate among multiple competing hypotheses. Some tests may involve passive inspection using sensors (e.g., use vision to determine if an access panel fastener is missing). Others might involve more active procedures which could compound the error if applied without due consideration (e.g., to check if the effector is jammed, rotating it in a confined area could damage nearby satellite components).

- **Diagnostic techniques.**

The capability to employ a variety of different styles of diagnosis as required by the problem at hand. Sometimes heuristic techniques will be applicable and serve to quickly localize a problem. In other cases, the ability to reason in depth about the structure and function of robot components will be important. The ability to use causal modelling techniques to project the effects of possible failures will be useful in other circumstances.

- **Explanation.**

The ability to describe and justify diagnostic conclusions to the level of detail as required by a system operator (see **Human Interface** below) or other knowledge based system. For effective error recovery, a diagnosis must have enough precision to specify the requirements for error recovery (e.g., "The forearm is broken" lacks sufficient detail for recovery). In addition, the ability to justify the particular conclusion reached (especially when there are closely competing alternative hypotheses) will be important for human evaluation and risk assessment of possible recovery procedures.

B.5 ERROR RECOVERY PLANNING

After a diagnosis has been determined, it is the role of error recovery procedures to plan how to get the plan back on track. This may involve small modifications to the plan at hand, the replanning of significant procedures, or even the abandonment of goals deemed impossible to accomplish in the changed circumstances. Some capabilities which will be important are:

- **Selection of predefined procedures.**

In some contingencies, recovery procedures will have been devised beforehand. The robot should have the ability to determine the applicability of these recovery procedures from a machine or human generated diagnosis.

- **Discrepancy analysis.**

The ability to determine what the effects of a particular failure are on planned actions, scheduled events, or actions in progress. In major failures, much of the plan remaining to execute will become invalid. However, in other cases the effects of failures may be slight and the existing plan resumed after some simple local error recovery. Discrepancy analysis will be critical in not wasting the work already put into planning and plan execution.

- **Ad hoc planning and plan integration.**

If predefined recovery procedures are inapplicable to the situation at hand, the robot must have the ability to specify the requirements of a novel procedure, invoke the appropriate planning process, and then incorporate the new procedure into the existing plan. Precision in the ability to excise the invalid components of the old plan and *splice in* the new procedure will be probably depend on **Discrepancy Analysis** as described above.

- **Goal maintenance.**

The robot must be able to determine when a plan is irrecoverable and a goal must be abandoned or revised. In addition to the case where recovery is

impossible, in some cases recovery may be too time consuming, too risky for the robot or workpiece, or unsafe for astronauts working nearby (e.g., venting fuel from a clogged or kinked refueling hose).

- **Revising and acquiring knowledge.**
If diagnosis processes have determined that some aspect of the robot's knowledge was in error, misapplied, or missing, a recovery procedure must determine how to correct the problem. This is one aspect of the machine learning problem, resolution of which will certainly be important in the achievement of full robot autonomy.

B.6 SIMULATION AND PREDICTION

Central to much of the robot's ability to solve problems is the capability to create and reason about the potential extended effects of alternative actions or events. In order to plan, the cumulative effects of alternative actions must be considered. In order to anticipate the long-term effects of potential failures, the immediate effects of the failure must be extended forward through time. Sometimes, in order to discover the cause of current anomalies, the ability to reason how the effects of hypothetical *previous* failures could propagate will be important.

- **Causal Simulation.**
The ability to utilize causal models to envision the effects or causes of particular states or events which concern the robot.

B.7 REAL-TIME PROBLEM SOLVING AND CONTROL

Of considerable concern is the robot's ability to behave in real-time. Situations which are beyond the robot's ability to "freeze" will be common (e.g., moving obstacles) and high performance will be required. The following capabilities are important:

- **Meta-level control.**
The ability to utilize knowledge about on-going problem solving to judge the relative importance, efficiency, reliability, and potential for success of alternative problem-solving strategies which are competing for real-time computing resources. The ability to make a quick and accurate judgement about where to focus problem-solving will be important in achieving high performance.

B.8 INTEGRATION

The Telerobot requires that a diverse set of software and hardware must smoothly integrated. From the point of view of the AI systems involved, two capabilities are critical:

- **Integration of multiple knowledge based systems.**
Much of this paper describes capabilities which will be provided by knowledge based systems. The ability to integrate this eclectic group of systems into a single functional unit is critical in achieving the full scope of robot autonomy.

In the past, the technical problems associated with each of the areas described have generally been studied in relative isolation from one another. A fully autonomous system, however, must rely on multiple knowledge based systems which can interact, cooperate, or merely tolerate one another in a common computing environment.

- **Integration with non-ai systems.**
The knowledge based systems, part of the robot's task planning and reasoning component, must be able to interact with the other non-AI systems which compose the bulk of the robot's control, sensory, and human interface abilities. Protocols for interaction with these systems must be established.

B.9 UNCERTAINTY MANAGEMENT

As the System Autonomy research plan states, uncertainty management is "...the ability to make sensible judgements and carry out reasonable actions when world knowledge is imprecise or incomplete, heuristics or models have built-in uncertainty, or actions have uncertain effects." Any system which behaves intelligently and robustly in the real world must account for the inherent lack of precision in knowledge and ability to control the world. The following capabilities will be important:

- **Identification of sources of uncertainty.** The robot must be able to locate and accommodate a wide variety of sources of uncertainty. For example, uncertainty can arise from partial or imprecise models, such as gaps in knowledge about the cause and effects of events or properties of objects (including the robot itself), lack of knowledge about the effects of active physical processes, or imprecision in modelling the intentions and plans of other agents.

B.10 HUMAN INTERFACE AND INTERACTION

Several human operator responsibilities must be supported by the human interface to the Artificial Intelligence components of the Telerobot:

- **Supervision.**
The operator must have the ability to determine, select, and specify goals for the robot to achieve prior to and during plan execution. The operator must have the ability to preview and select among alternative plans suggested by the robot. At all times during plan execution, the operator must have the ability to suspend or redirect the activities of the robot.
- **Criticism.**
The operator must have the ability to evaluate and criticize plans. This could include the ability to make modifications to plans (e.g., manipulator trajectories) which are otherwise acceptable to the robot.
- **Cooperation.**
The operator must have the ability to be involved in both the problem-solving activities of the robot as well as the actual execution of plans themselves. Especially in the early stages of development, the robot will often lack critical

knowledge about how to locate or recover from specific errors which occur during plan execution. The human operator must have the ability to instruct the robot and/or carry out the operations alone. In later Telerobot development, the ability for the robot and human to coordinate their activities will be important

B.11 KNOWLEDGE ACQUISITION AND LEARNING

There is a difficult bottleneck in the creation of knowledge based systems, i.e., the identification, acquisition, representation, verification, and management of the knowledge which is required for problem solving. Knowledge-engineering has today developed into a skilled craft with numerous tools to aid a human developer of knowledge based systems. However, the volume and complexity of the knowledge required for the Telerobot (and other complicated systems) is sure to overwhelm the techniques which exist today. Ultimately, the robot should be able to perform much of the knowledge acquisition and maintenance problem autonomously, i.e., learn. Machine learning has the role of resolving uncertainty, correcting knowledge errors or gaps, and in generating new capabilities for the robot. The following capabilities leading to autonomy are important:

- **Use of CAD/CAM databases.**

Effective diagnosis, assembly, inspection, and other tasks designated for the robot will require detailed knowledge about the structure and function of the objects it manipulates, such as satellites. Much of this information is expected to be available in computable form in CAD/CAM databases. Techniques must be developed for exploiting this information and transforming it into representations usable for knowledge based problem-solving.

- **Use of human documentation.**

Many of the tasks which the Telerobot is expected to automate have been designed for humans. To a certain extent, human readable documentation exists on these tasks or the systems to be manipulated. Techniques for utilizing this knowledge resource should be developed, including natural language parsing.

- **Management of massive knowledge bases.**

There are numerous problems associated with managing the volume of knowledge which will be required by the robot. There will probably be multiple representations of the same information which must be consistent. At any given stage of problem-solving, only a subset of knowledge is required. This must be quickly and efficiently provided. The knowledge base management problem appears to subsume many traditional database management problems.

- **Learning by experience.**

When the robot makes a mistake during plan execution, and identifies the problem as imprecision in its knowledge, it shouldn't make the same mistake a second time. In some cases, the knowledge is correct, but simply inapplicable to the current situation in which the error occurred. In this case, the robot must be able to recognize in future situations when the same mistake could occur, recall the correction it devised previously, and implement the correction in the current situation.

- **Learning by discovery.**
The robot should have the ability to fortuitously notice or bring about situations which are instructive.
- **Learning by teaching.**
The robot should have the ability to acquired knowledge from direct interaction with humans, either through factual presentation, reasoning from examples, or other methods of instruction.

C. TECHNICAL AREAS FOR RESEARCH

This section presents some of the technical areas for Artificial Intelligence research which should be supported by the Telerobot project over the next several years to achieve some of the capabilities noted above. NASA probably does not need to sponsor work to achieve all of these capabilities; many required capabilities cross application boundries and the necessary research is currently funded by other government agencies. NASA does need, however, to focus research on those areas which are critical for the near-term establishment and success of the Telerobot (circa 1993). The following areas are relevant:

- **Planning and scheduling.**
Near term: Research on conditional, contingency, and least commitment planning. Research on incorporation of physical, spatial, and temporal constraints and associated planning processes with task activity planning. Research on resource management. Management of uncertainty in planning.
Longer term: Research on multi-agent cooperation, including: Plan recognition, communication of plans and intentions between agents, command and information request communications, real-time compensation for other agent action discrepancies, supervisory versus distributed control issues.
- **Spatial planning and reasoning.**
Near term: Spatial model representations and databases which are useful for problem-solving tasks and easily integrable with knowledge based systems. Qualitative causal modelling of physical processes. Reasoning about spatial and physical constraints on task planning processes. Manipulator trajectory planning and collision avoidance. *Longer term:* 3-D robot mobility and navigation.
- **Execution Monitoring.**
Near term: Sensor planning issues, including: Plan action analysis to support sensor allocation, sensor resource allocation and real-time management, formal languages for specification of sensor plans, issues in active versus passive sensing and plan interactions, selective versus exhaustive monitoring issues, grain-size of monitoring, expectation generation. Monitoring issues such as: symbolic classification of sensor data, partial matching. *Longer term:* Issues in multi-sensor and temporal data fusion, uncertainty management, partial matching, noticing unexpected events.
- **Diagnosis.**
Near term: Identification of source of plan failures, including system failures, knowledge failures, or unexpected world states or effects of actions. Utilization

of multiple kinds of diagnostic knowledge, including heuristic, first principles, procedural. *Longer term:* Classification of novel failures, localizing multiple failures, assumption changes, troubleshooting knowledge failures.

- **Error recovery.**

Near term: Analysis of discrepancies on existing plan, plan representation and operators for combining plan fragments such as recovery plans. Integration with planning processes. *Longer term:* Determining discrepancy effects on planning knowledge (a learning problem), including: recognizing that knowledge is at fault, determining which knowledge, explaining why it is wrong, and implementing corrective action.

- **Simulation and prediction.**

Near term: Causal modelling for the Telerobot domain, issues in modelling temporal constraints and physical processes. *Longer term:* Dealing with bad or ambiguous causal models.

- **Real-time problem-solving.**

Near term: Issues in goal, planning, and execution management, including: goal valuation and priority, goal viability, recognition and management of dynamic resources required to achieve goals, interaction with human goals. *Longer term:* Issues in parallel or distributed processing of knowledge based systems. Real-time problem-solving software architectures.

- **Integration.**

Near term: Issues in integration of multiple knowledge based systems, including: Message and request communication, shared knowledge, executive versus distributed control. Issues in integration with non-ai systems, including superordinate and subordinate roles. *Longer term:* Integration architectures, including blackboards, distributed computing systems.

- **Human interface.**

Near term: Knowledge based system command languages, graphical and iconic display of plans, interface for supervised plan execution. *Longer term:* Natural language understanding and generation, mixed initiative dialogs, generation of explanations of knowledge based system behavior, communication of shared knowledge.

D. METHODOLOGY

The Telerobot project involves research and development in Artificial Intelligence and other areas which spans the range from highly risky and innovative basic research through applied research and engineering to application engineering of operational systems. It is essential that a flow of information, techniques, skills, and even personnel be maintained across this span of research and development. A recent perspective article in *Science* by Dr. A. M. Clogston of AT&T Bell Laboratories may be instructive:

"...The research-development interface is a difficult enough barrier to surmount, even within a highly integrated R&D laboratory, and it is

more difficult to import research, even with the best will on both sides. The best way to import university research into an R&D laboratory is through active in-house basic and applied research groups, a channel that would be closed if the laboratory relied too much on external research."

- A. M. Clogston, "Applied Research: Key to Innovation", *Science*, 235, 4784, (1987)

Our methodological goal should be to combine a set of tightly focussed external research contracts with substantive in-house basic and applied research directed at:

- Importing this technology and technology sponsored by other agencies into the Telerobot project.
- Filling in the "gaps" in basic and applied research areas which are important for Telerobot.
- Performing the applied research necessary to move conceptual breakthroughs into operational demonstrations of robot capability.

To facilitate the importing of university technology, individual graduate students and professors conducting research sponsored by, or of interest to, the Telerobot project will be invited to visit the laboratory and work with JPL personnel for short periods of time. This should enhance the education of JPL personnel in emerging AI technology as well as provide the necessary project visibility and feedback to basic research efforts.

E. CONCLUSIONS

~~This paper has presented~~ a view of the capabilities and areas of artificial intelligence research which are required for autonomous space telerobotics extending through the year 2000. In the coming years, JPL will be conducting directed research to achieve these capabilities, as well as drawing heavily on collaborative efforts conducted with other research laboratories.

F. ACKNOWLEDGEMENT

The research described in this paper was done at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.